Exact rule-based stochastic simulations for the system with unlimited number of molecular species

Anton V. Bernatskiy[1, a)] and Elizaveta A. Guseva[2, b)]
[1)] *Center for Complex Systems, University of Vermont, 210 Colchester Ave, Burlington, VT, 05405*
[2)] *Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, NY 17794-5252*

(Dated: 5 October 2016)

We introduce expandable partial propensity direct method (EPDM) – a new exact stochastic simulation algorithm suitable for systems involving many interacting molecular species. The algorithm is especially efficient for sparsely populated systems, where the number of species that may potentially be generated is much greater than the number of species actually present in the system at any given time. The number of operations per reaction scales linearly with the number of species, but only those which have one or more molecules. To achieve this kind of performance we are employing a data structure which allows to add and remove species and their interactions on the fly. When a new specie is added, its interactions with every other specie are generated dynamically by a set of user-defined rules. By removing the records involving the species with zero molecules, we keep the number of species as low as possible. This enables simulations of systems for which listing all species is not practical. The algorithm is based on partial propensities direct method (PDM) by Ramaswamy et al. for sampling trajectories of the chemical master equation.

arXiv:1609.06403v2 [cs.DS] 4 Oct 2016

Frequently used abbreviations and symbols

SSA – Stochastic Simulation Algorithm
DM – Direct Method[17]
PDM – Partial propensity Direct Method[34]
SPDM – Sorting Partial propensity Direct Method[34]
EPDM – Expandable Partial propensity Direct Method
N – number of molecular or agent species
M – number of all possible reactions or agent interactions
α – maximum number of possible reactions or agent interactions between any pair of molecular or agent species

I. INTRODUCTION

Mathematical modeling of chemical reactions is an important part of systems biology research.

The traditional approach to modeling chemical systems is based on the law of mass of action equations which assumes reactions to be macroscopic, continuous and deterministic. It doesn't account for the discreteness of the reactions or temporal and spatial fluctuations, describing only the average properties instead. This makes it ill-suited for modeling nonlinear systems or systems with a small number of participating molecules, such as living cells.

In contrast, *stochastic simulation algorithms* (SSAs) do account for discreteness and inevitable randomness of the process. As a result, these methods are becoming increasingly popular in modern theoretical cell biology[1,24,40]. Additionally, they have more physical rigor compared to the empirical law of mass action ordinary differential equations[18].

It is worth mentioning that replacing molecules with any other interacting agents does not invalidate the approach, as long as certain conditions are met. This makes it suitable for modeling many non-chemical systems in areas such as population ecology[29,37], evolution theory[13,26], immunology[11], epidemiology[5,25,27], sociology[12,39], game theory[30], economics[23,31], robotics[3,4] and information technology[22].

The majority of contemporary SSAs are based on the algorithm by Gillespie[16–18] known as *direct method* (DM). It describes a system with a finite state space undergoing a continuous-time Markov process. The time spent in every state is distributed exponentially; future behavior of the system depends only on its current state. Gillespie has shown[18] that for chemical systems these assumptions hold if the system is well-stirred and molecular velocities follow Maxwell-Boltzmann distribution. Because of DM's good physical basis its solutions are as accurate or more accurate than the law of mass equations, which isn't necessarily a good predictor of mean values of molecular populations[17].

Time complexity of DM is linear in the number of reactions. For highly coupled systems that means that the run-time grows as a square of the number of species. This poses a problem for many models in systems biology which describe systems with a multitude of molecular species connected by complex interaction networks. Storage complexity of the DM is linear in the number of possible reactions.

Many alternative SSAs which improve the time complexity of DM (see Table I for an incomplete list) were developed. Approximate SSAs make such improvements at the cost of introducing additional approximately sat-

[a)] Electronic mail: abernats@uvm.edu
[b)] Electronic mail: eguseva@tutanota.com

isfied assumptions, while exact SSAs achieve better runtime by performing a faster computation that is equivalent to the one performed by DM. Exact methods have been developed with time per reaction that is linear[34] or, for sparsely connected networks, even constant[35] in the number of molecular species.

One deficiency shared by most SSAs is that they operate with a static list of all possible reactions and species, while for some systems of interest maintaining such a list is not possible. An example of such system is heterogeneous polymers undergoing random polymerization, an object of interest for the researchers of prebiotic polymerization and early evolutionary processes. The number of heteropolymers that can be produced by such a process is infinite. Even if we only consider species of length up to L, we have to deal with $\mathcal{O}(p^L)$ species, where p is the number of monomer species. This causes both time and storage complexity of stochastic models to grow exponentially with L, which limits the studies to very short polymers.

However, if an algorithm can maintain a dynamic list of molecular species and interactions, the complexity of modeling such systems can be substantially reduced. In our example, the set of possible species is so large that even for moderate cutoff L the vast majority of all possible species will have a population of zero. If at any time step a specie has a population of zero, no reaction can happen in which this specie is a reagent. Therefore, the reactions involving such a specie need not be tracked until some reaction occurs in which the specie is a product.

In addition, many studies are concerned with emergent phenomena, which often involve the system exhibiting some dynamical pattern, e.g. sitting in a stable or chaotic attractor. Such phenomena tend to only involve a subset of possible states and transitions and not explore the space of all possibilities uniformly. This may limit the species and reactions involved in the phenomenon to a small subset of all possible species and reactions. For example, in the study by Guseva et al.[20] effective populations of all sequences are $\propto 10^3$, while the number of all the possible species is $\propto 10^7$. The model is tightly coupled, thus time costs are 10^4 times higher and memory costs are 10^8 times higher than they could be if only the final subset of reactions and species was considered and the SSA retained the time and storage complexity of the best SSAs for static lists (e.g. PDM[34]).

However, not only this subset is unknown prior to the experiment, but the transient to the behavior of interest may involve many more species and reactions than the behavior itself. Thus, the simulation cannot be constrained to use a smaller list of species and reactions if such list is static.

One class of systems in which reactions cannot be listed occurs in solid state chemistry. To simulate those, Henkelman and Jonsson[21] developed an algorithm in which the list of possible reactions is generated on the fly and used in the standard Gillespie's algorithm. This approach was formulated as a part of simulation framework specific to solid state chemistry. It remained obscure outside of the solid state chemistry community and was rediscovered independently by authors of the present work early in the course of its preparation.

Here we present *extendable partial-propensity direct method* (EPDM): a general purpose, exact SSA in which species and reactions can be added and removed on the fly. Unlike Henkelman and Jonsson's approach, our algorithm is based on the *partial-propensity direct method* (PDM) by Ramaswamy et al.[34] and retains its linear time complexity in the number of species. Storage complexity is linear requirements in the number of reactions.

Only the species with nonzero populations and the reactions involving them are recorded and factored into the complexity. This reduces time complexity of executing one reaction by a factor of N_{tot}/N where N_{tot} is the number of all possible species and N is the number of species with nonzero population at the time of the reaction. Storage complexity is reduced by the square of that factor for densely connected reaction networks.

We test the performance of a C++ implementation of our algorithm[1] for two chemical systems, investigate scaling and compare the performance with two other SSAs. The results confirm our predictions regarding the algorithm's complexity.

Exact	• Direct method (DM)[16,17] • First reaction method (FRM) [16] • Gibson-Bruick's next-reaction method (NRM)[15] • Optimized direct method (ODM)[10] • Sorting direct method (SDM)[28] • Partial propensity direct method (PDM)[34]
Approximate	• τ-leaping[7,9,19,32] • k_α-leaping[19] • Implicit τ-leaping[36] • The slow-scale method[8] • R-leaping[2] • L-leap[33] • K-leap[6]

TABLE I. List of some exact and approximate stochastic simulation algorithms

II. DIRECT STOCHASTIC ALGORITHMS NOW

A. Gillespie Algorithm

Being an exact SSA, our algorithm performs an optimized version of the same basic computation as the Gillespie's DM[14,17]. DM is based on the following observation:

[1] https://github.com/abernatskiy/epdm

> Probability that any particular interaction of agents occurs within a small period of time is determined by and proportional to a product of a rate constant specific to this interaction and a number of distinct combinations of agents required for the interaction.

Since we developed EPDM with chemical applications in minds, we will hereafter refer to agents of all kinds as molecules. However, as long as the observation above holds, the method is valid for any other kind of agent (see Introduction).

The algorithm starts with initialization (**Step 1**) of the types and numbers of all the molecules initially present in the system, reaction rates and the random number generator. Then propensities of all reactions are computed (**Step 2**). Propensity a_μ of a reaction R_μ ($\mu \in \{1, 2, ..., M\}$) is proportional to its reaction rate c_μ:

$$a_\mu = h_\mu c_\mu. \tag{1}$$

Here, h_μ is the number of distinct molecular reactant combinations for reaction R_μ at the current time step, a combinatorial function which depends on the reaction type and the numbers of molecules of all reactant types[16]. Total propensity is the sum of propensities of all reactions:

$$a = \sum_{\mu=1}^{M} a_\mu. \tag{2}$$

The next step (**Step 3**), called Monte Carlo step or sampling step, is the source of stochasticity. Two real-valued, uniformly distributed random numbers from $[0, 1)$ are generated. The first (r_1) is used to compute time to next reaction τ:

$$\tau = \frac{1}{a} \ln\left(\frac{1}{r_1}\right). \tag{3}$$

The second one (r_2) determines which reaction occurs during the next time step τ. The j-th reaction occurs if

$$\sum_{\mu=1}^{j-1} a_j \leqslant a r_2 < \sum_{\mu=1}^{j} a_j. \tag{4}$$

The next step is update (**Step 4**): simulation time is increased by τ generated at Step 3, molecules counts are updated using the stoichiometric numbers of the sampled reaction and propensities are updated in accordance with the new molecular counts.

The last step is iteration: go back to Step 3 unless some termination condition is met. Termination should occur if no further reactions are possible (i.e. when the total propensity $a = 0$). Optional termination conditions may include reaching a certain simulation time, performing a given number of reactions, reaching some steady state etc.

At every step the algorithm looks through the list of all M possible reactions. Therefore, the time it takes to process one reaction (i.e. perform Steps 3 and 4) is proportional to M:

$$t_{reac} = \mathcal{O}(M). \tag{5}$$

There must be a record for every reaction, so the space complexity is also $\mathcal{O}(M)$.

B. Partial propensity methods

For many systems it is valid to neglect reactions which involve more than two molecules or agents. This premise allows for a class of partial-propensity direct methods. The method described in the present work is among those.

If α is the maximum number of reactions which may happen between any pair of the reagent species then $M = \mathcal{O}(\alpha N^2)$. Then the expression for the time complexity of DM (5) becomes quadratic in N:

$$t_{reac} = \mathcal{O}(\alpha N^2). \tag{6}$$

Partial-propensity direct method (PDM) and sorting partial-propensity direct method (SPDM)[34] improve this bound to $\mathcal{O}(\alpha N)$ by *associating each reaction with one of the involved reagents* and *sampling the reactions in two stages*. In the first stage the first reactant specie of the reaction to occur is determined; this takes $\mathcal{O}(N)$ operations. The second stage determines the second specie and a particular reaction to occur; this takes $\mathcal{O}(\alpha N)$ operations, and the total complexity of the sampling step adds up to $\mathcal{O}(\alpha N)$.

Since any specie can be involved in at most $\mathcal{O}(\alpha N)$ unimolecular or bi-molecular reactions, only $\mathcal{O}(\alpha N)$ values have to be updated when the molecular counts change. This enables PDM and SPDM to perform the update step without worsening the time complexity of the sampling step.

The final time complexity of these algorithms

$$t_{reac} = \mathcal{O}(\alpha N) \tag{7}$$

holds irrespective of the degree of coupling of the reaction network. The number of records required for sampling is still $\mathcal{O}(M) = \mathcal{O}(\alpha N^2)$.

Note that for sparsely coupled reaction networks, time complexity can be improved further to $\mathcal{O}(1)$[35]; here we won't concern ourselves with such reaction networks.

III. ALGORITHM DESCRIPTION

A. Reaction grouping

Similar to PDM and SPDM[34], our method relies on splitting the reactions into groups associated with one

of the reagents. To simplify this procedure, we reformulate all reactions with up to two reagents as bimolecular by introducing the virtual *void specie* $\varnothing$ which always has a molecular count of 1. It can interact with other, real species and itself, however its stoichiometry is always zero. With these new assumptions, unimolecular reactions are reformulated as follows:

$$\begin{aligned} &S_i \to \text{products} \Rightarrow \\ &1\,S_i + 0\,\varnothing \to \text{products}. \end{aligned} \tag{8}$$

Source reactions are also reformulated:

$$\begin{aligned} &\varnothing \to \text{products} \Rightarrow \\ &0\,\varnothing + 0\,\varnothing \to \text{products}. \end{aligned} \tag{9}$$

Our algorithm keeps a list of species existing in the system to which entries can be added. When a new specie S_i is added to the list, the algorithm considers every specie S_j in the updated list. For each (unordered) pairing $\{S_i, S_j\}$ a list of possible reactions is generated. If the list is not empty, it is associated with S_j, the specie that has been added to the list earlier unless it coincides with S_i. If $S_i = S_j$, the list is associated with S_i.

For example, the first step of the initialization stage involves adding the first element to the list of species which is always the void specie $\varnothing$. At this point there are no species in the list aside from $\varnothing$, so the algorithm checks which kinds of reactions may happen between $\varnothing$ and itself. Due to the reformulation (9) this will involve all source reactions. Their list will be generated and associated with the newly added specie $\varnothing$.

Another example. Suppose the list of known species is $[\varnothing, S_1]$ and we're adding a specie S_2 which reacts with S_1, itself and also participates in some unimolecular reactions. The list becomes $[\varnothing, S_1, S_2]$ and the algorithm proceeds to pair up S_2 with every specie in it and generate reaction lists.

Due to the reformulation (8) it will find all the unimolecular reactions for the pair $\{S_2, \varnothing\}$. The list of these reaction will be associated with the previously known specie of the pair, $\varnothing$.

When considering the pair $\{S_2, S_1\}$ it will find all the reactions between S_1 and S_2 and associate them with S_1.

Finally, it will consider a pair $\{S_2, S_2\}$ and find its reactions with itself. Since there are no previously known species in the pair, the reactions will be associated with S_2.

B. Propensities generation

As the reactions are generated and associated with species we also compute and store some propensities.

For all reactions R_μ with up to two participating molecules full propensities a_μ are

$$\begin{aligned} &a_\mu = n_i n_j c_\mu \text{ for a bimolecular reaction} \\ &\quad \text{of distinct species } i \neq j,\ S_i + S_j \to \text{products}, \\ &a_\mu = \frac{1}{2} n_i (n_i - 1) c_\mu \text{ for a bimolecular reaction} \\ &\quad \text{of identical species } 2S_i \to \text{products}, \\ &a_\mu = n_i c_\mu \text{ for a unimolecular reaction} \\ &\qquad\qquad S_i \to \text{products}, \\ &a_\mu = c_\mu \text{ for a source reaction } \varnothing \to \text{products}. \end{aligned} \tag{10}$$

Since the void specie $\varnothing$ has a fixed population of 1 we can use the formula for the unimolecular reactions for source reactions as well: $a_\mu = 1 \cdot c\mu$ for $\varnothing \to$ products. Then for all reactions we can define *partial propensity w.r.t. the reactant specie* S_i as

$$\pi_\mu^{(i)} \equiv a_\mu / n_i. \tag{11}$$

For the reactions with up to two reagents, partial propensities are[34]

$$\begin{aligned} &\pi_\mu^{(i)} = n_j c_\mu \text{ for a bimolecular reaction} \\ &\quad \text{of distinct species } i \neq j,\ S_i + S_j \to \text{products}, \\ &\pi_\mu^{(i)} = \frac{1}{2}(n_i - 1) c_\mu \text{ for a bimolecular reaction} \\ &\quad \text{of identical species } 2S_i \to \text{products}, \\ &\pi_\mu^{(i)} = c_\mu \text{ for a unimolecular reaction} \\ &\qquad\qquad S_i \to \text{products}, \\ &\pi_\mu^{(\varnothing)} = c_\mu \text{ for a source reaction } \varnothing \to \text{products}. \end{aligned} \tag{12}$$

Suppose some specie S_j is added to the list of known species as described in section III A. For every specie S_i in the updated list we generate a list of possible reactions between S_i and S_j and associate it with S_i. For every reaction R_{ijk} in the list, we compute its rate constant c_{ijk} and its partial propensity $\pi_{ijk}^{(i)}$ w.r.t. S_i using formulas (12).

We also keep some sums of propensities to facilitate the sampling. Given a list of reactions between S_i and S_j associated with S_i, we define $\Psi_{ij}^{(i)}$ as

$$\Psi_{ij}^{(i)} = \sum_{k=1}^{\alpha_{ij}} \pi_{ijk}^{(i)}, \tag{13}$$

where α_{ij} is the number of reactions possible between S_i and S_j.

$\Lambda_i^{(i)}$ is the sum of partial propensities of all reactions associated with S_i:

$$\Lambda_i^{(i)} = \sum_{j=1}^{m_i} \Psi_{ij}^{(i)}. \tag{14}$$

Here, m_i is the number of lists of reactions with other species associated with S_i.

Σ_i is the full propensity of all reactions associated with S_i:

$$\Sigma_i = n_i \Lambda_i^{(i)}. \tag{15}$$

and a is defined as a total full propensity of all reactions in the system:

$$a = \sum_{i=1}^{N} \Sigma_i. \tag{16}$$

C. Data model

Similar to Steps 3 and 4 in Gillespie's DM (see section II A), execution of each reaction in our algorithm involves two steps: *sampling* and *updating*. We'll refer to the data used in the sampling process as *primary* and call all the data which is used only for updating *auxiliary*.

To minimize the overhead associated with adding and removing the data we use a hierarchical linked list (multi-level linked list of linked lists). We define several data structure types (see table II), many of which include lists of instances of other types. Complete data model is shown in Figure 1.

Top level list is stored within a structure of type `TotalPopulation`. Aside from the list, the structure holds the data describing the system as a whole: total propensity a (see (16)), current simulation time t and the random generator state.

Elements of the top level list are structures of type `Population`, each of which holds the information related to the molecular population of a particular specie S_i. The information about the specie itself, to which we will in the context of its `Population` refer as the *owner specie*, is represented by a member object of user-defined class `Specie` (see note 2 at the end of this section), containing a unique string specie identifier ID_i as a member variable. `Population` also contains the number of molecules in the population n_i, total propensity of the population Σ_i and total partial propensity $\Lambda_i^{(i)}$ w.r.t. S_i of all reactions associated with S_i (see eqs. (15) and (14)).

The `Population` structures are appended to the top level list sequentially over the course of the algorithm execution (see section III A). Each of them holds, in addition to the data mentioned above, two linked lists: of structures of type `Relation` and of structures of type `RelationAddress`. The former stores the list of all reactions possible between the owner specie and some other specie that has been added later in the course of the algorithm's operation. `RelationAddress` structures hold the references used to access `Relation`s in which the owner specie participates, but not as an owner. Those include all the species added before the owner.

It can be observed that the `Population` which has been added first will necessarily has an empty list of `RelationAddress`es and a potentially large list of `Relation`s, with as many elements as there are species with which the first specie has any reactions. On the other hand, the `Population` added last will not be an owner of any `Relation`s and its list of those will be empty. In this case, all information about the specie's reactions is owned by other species and only available within the `Population` through the list of `RelationAddress`es.

A `Relation` structure owned by S_i holds a linked list of all possible `Reaction`s between S_i and S_j and their total partial propensity $\Psi_{ij}^{(i)}$ (see eq. (13)). `Reaction`s store the reaction rate c_{ijk}, partial propensity $\pi_{ijk}^{(i)}$, a table of stoichiometric coefficients and IDs of all participating species, including products. Additionally, each of them stores an auxiliary reference to the `RelationAddress` structure pointing to this `Relation`.

`RelationAddress` is an auxiliary structure holding references to a `Relation` in which the owner specie participates without owning it, and to the `Population` of the specie which does own the `Relation`. It also contains the references to itself and to the list holding it.

Notes:

1. We will say that a `TotalPopulation` is *valid* iff all variables mentioned in eqs. (12–16) satisfy these equations and all references mentioned in the present section are valid and point where intended.

2. Unique specie string ID format and `Specie` objects are user-defined and must be freely convertible between each other. The utility of having a separate `Specie` class lies in it having a user-defined method `reactions()` : $(self, \texttt{Specie}) \rightarrow ListOfReactions$ which produces a list of all reactions possible between the specie described by a caller object and a specie described by the the argument object. For some systems, this computation can be made much faster if some auxiliary data can be kept in the structure describing a specie. String IDs on the other hand are intended as memory efficient representations of specie data which are used when this computation may not be needed, e.g. to represent reaction product not yet present in the system or for specie comparison.

   Our method does not rely on this separation other than as a means of optimization.

3. Additionally, in our implementation we provide a global structure to hold the parameters of the system which may influence the behavior of the `reactions()` method of class `Specie`.

D. Adding a `Population`

Suppose that we have a valid `TotalPopulation` with $N-1$ `Population`s. To add a `Population` based on a specie ID and a molecular count n, we follow the method described in section III A (see also algorithm 1). We begin by appending a `Population` p_N to

Data type	Necessary members				Significance
	Containers	Scalars	References	Functions	
`TotalPopulation`	Linked list of `Population`s	Total propensity a (16), current time t, random generator state	–	–	Represents the whole system being modeled
`Population`	Linked list of `Relation`s, linked list of `RelationAddresse`s	`Specie` S, molecular count n, propensity sums Σ (15) and Λ (14)	–	–	Represents a population of molecules of a specie S
`Relation`	Linked list of `Reaction`s	Partial propensity sum Ψ (13)	To `RelationAddress` pointing to this `Relation`	–	Data on all reactions possible between a pair of species, to be stored withing the list of an *owner* specie
`RelationAddress`	–	–	To a `Relation`, its owner's `Population`, list containing this `RelationAddress` and itself	–	Reference to a `Relation` to be kept by a non-owner specie, useful in propensity updates and population deletions
`Reaction`	Array of pairs (ID, σ) of IDs and stoichiometries of participating species	Rate constant c, partial propensity π w.r.t. the owner specie (12)	–	–	Represents a reaction
`Specie`	–	Unique, compact specie identifier ID	–	`Specie::reactions(Specie)`	An extended specie representation capable of keeping extra information to generate lists of possible reactions with `reactions()` method quickly

TABLE II. List of data structure types used by the algorithm

`TotalPopulation`'s list. p_N's `Specie` structure S_N is converted from ID, its molecular count $n_N \leftarrow n$ and propensities $\Lambda_N^{(N)} \leftarrow 0$, $\Sigma_N \leftarrow 0$.

Next, for each `Population` p_i in `TotalPopulation`'s updated list (including the newly added p_N) we compute a list of all possible reactions between S_N and p_i's `Specie` S_i. If the list is not empty, it is then converted into `Relation` ρ_{iN} and stored at p_i's list of those. To do the conversion, we compute partial propensities $\pi_{iNk}^{(i)}$ and $\Psi_{iN}^{(i)}$ using eqs. (12) and (13). We also store copies of S_i's ID at every `Reaction` in the list and at the `Relation` ρ_{iN} itself, to keep track of the specie w.r.t. which the current partial propensities are computed. ρ_{iN}'s reference to the `RelationAddress` structure is invalid at this point.

After appending the `Relation` we update p_i's propensities: $\Lambda_i^{(i)} \leftarrow \Lambda_i^{(i)} + \Psi_{iN}^{(i)}$, $\Sigma_i \leftarrow n_i \Lambda_i^{(i)}$.

We then proceed to create a `RelationAddress` structure RA_{Ni} pointing to ρ_{iN}. A blank structure is appended to p_N's list. We take references to ρ_{iN}, p_i, RA_{Ni} and p_N's list of `RelationAddresse`s and save them at RA_{Ni}.

Then, we take the reference to RA_{Ni} and store it at ρ_{iN}. At this point, all references in the whole structure are valid and correct.

Processing each p_i in the `TotalPopulation`'s list takes $\mathcal{O}(\alpha)$, so the whole procedure up to this point takes $\mathcal{O}(\alpha N)$. At this point we need to update the total propensity of the system a by recomputing it, which takes $\mathcal{O}(N)$[2].

The total number of operations it takes to add one `Population` is $\mathcal{O}(\alpha N)$. The operation preserves the validity of the data structure.

E. Initialization stage

To initialize the data structure, we make a `TotalPopulation` with an empty list of `Population`s, initialize $a \leftarrow 0$, $t \leftarrow 0$ and the random generator with a seed. The only variable that needs to take a particular value for the `TotalPopulation` to be valid is a and it has the correct value of 0; therefore, it is a valid `TotalPopulation`. We build the data structure by adding the initial populations to this structure as described in section III D. The resulting structure is valid because we only used validity-preserving operatins.

Adding every population takes $\mathcal{O}(\alpha N)$ operations and it must be repeated N times. This brings the complexity of the initialization step to $\mathcal{O}(\alpha N^2)$.

[2] It can be done during the iteration in $\mathcal{O}(1)$, but we chose to recompute it for improved numerical accuracy.

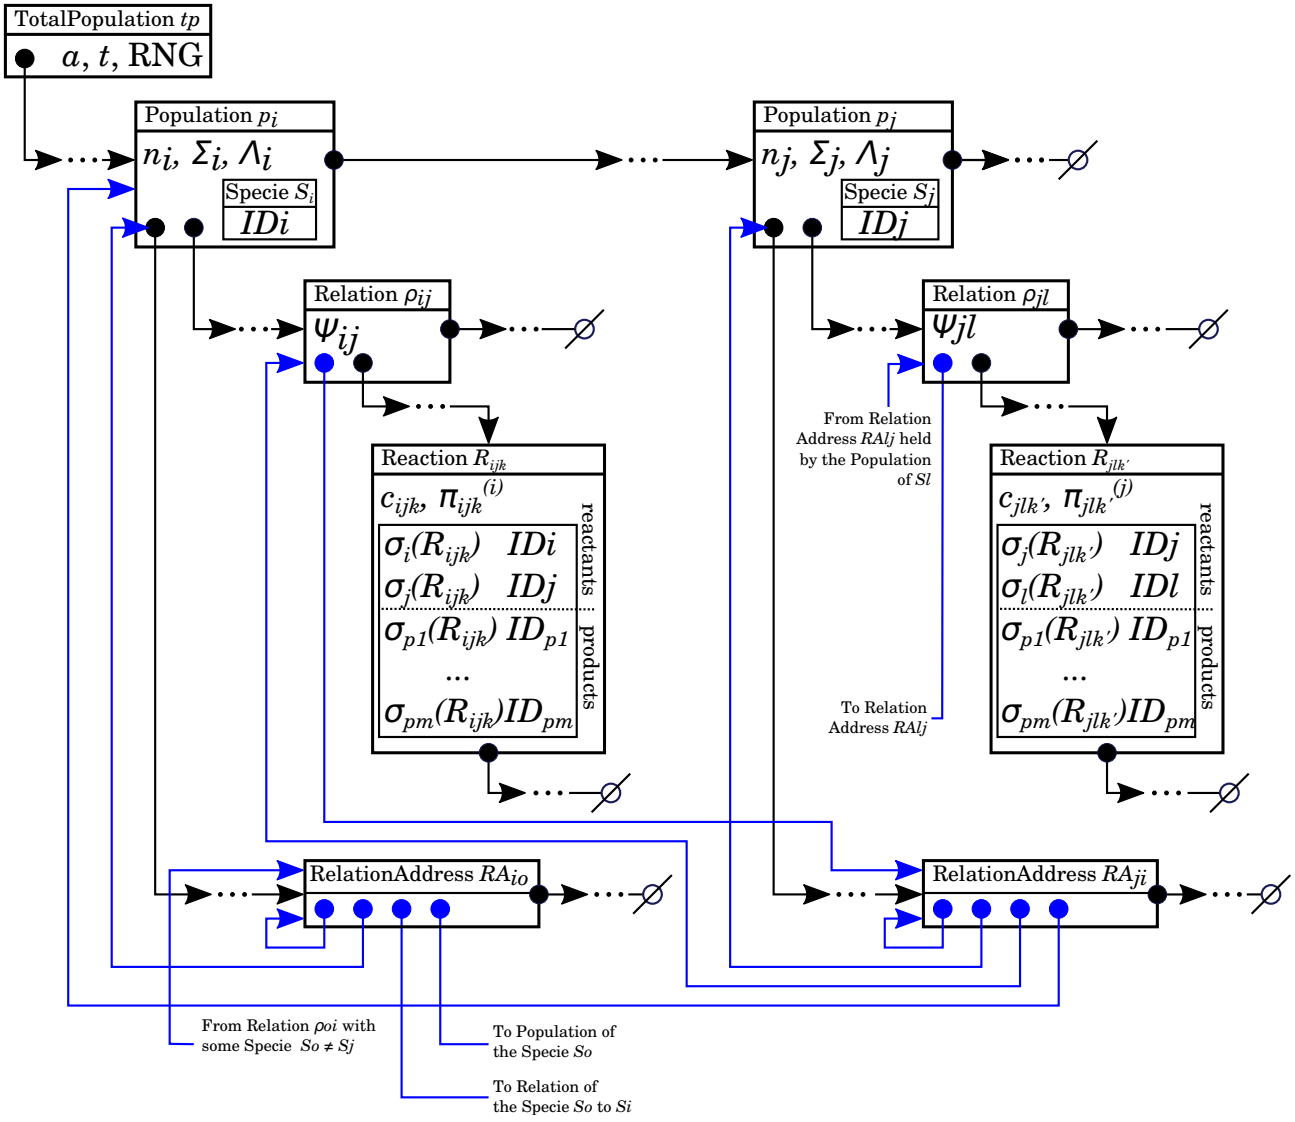

FIG. 1. EPDM data model. Circles with arrows denote references. References shown in black form the hierarchical linked list; in implementation, many of those are hidden within the `std::list` container. Auxiliary references added for dynamic updates are shown in blue. Letters and boxes within the boxes indicate non-reference member variables. Crossed empty circle represents a null reference.

F. Deleting a `Population`

Our algorithm is designed to keep track only of the species with a nonzero molecular count and reactions involving them. To accomplish that, we use the addition operations described in the previous section and deletion operation described here (see also algorithm 2). Deletion is only ever applied to `Populations` of species with zero molecules. All propensities of reactions involving such species are zeros; this enables us to simplify the procedure.

To remove a `Population` p_i, we begin by removing all `RelationAddress` structures pointing at `Relations` in p_i's list. Each `Relation` ρ_{ij} in p_i's list contains a reference to the `RelationAddress` structure pointing at it, RA_{ji}, which in turn contains a reference to itself and to a list holding it. We use those to remove each RA_{ji} from its list. Since S_i can be involved in at most N relations, this step takes $\mathcal{O}(N)$ operations.

Next, we remove all `Relations` in which S_i participates, but which are owned by other species. For each `RelationAddress` RA_{ij} we follow its references to the `Population` p_j of the other specie and to its `Relation` ρ_{ji} with S_i. We use those to remove ρ_{ji} from p_j's list. This step also takes $\mathcal{O}(N)$ operations.

Finally, we delete the whole `Population` structure p_i from `TotalPopulation`'s list. We recursively remove all the structures in its lists, which takes $\mathcal{O}(N)$ operations for the list of `RelationAddresses` and $\mathcal{O}(\alpha N)$ operations for the list of `Relations`. The resulting

Algorithm 1 Adding a `Population`

1: **function** AddPopulation(tp, ID, n)
2: Append a new `Population` p_N to tp's list
3: Convert ID to a `Specie` object S
4: $S_N \leftarrow S$, $n_N \leftarrow n$
5: $\Lambda_N^{(N)} \leftarrow 0$, $\Sigma_N \leftarrow 0$
6: **for all** `Population`s p_i in tp's list **do**
7: $ListOfReactions \leftarrow S_N$.`reactions`($S_i$)
8: **if** $ListOfReactions$ is empty **then**
9: **continue**
10: **end if**
11: Append a new `Relation` ρ_{iN} to p_i's list
12: Append all `Reaction`s in $ListOfReactions$ to ρ_{iN}'s list
13: $\Psi_{iN}^{(i)} \leftarrow 0$
14: **for all** `Reaction`s R_{iNk} in ρ_{iN}'s list **do**
15: Compute $\pi_{iNk}^{(i)}$ using (12)
16: $\Psi_{iN}^{(i)} \leftarrow \Psi_{iN}^{(i)} + \pi_{iNk}^{(i)}$
17: **end for**
18: $\Lambda_i^{(i)} \leftarrow \Lambda_i^{(i)} + \Psi_{iN}^{(i)}$
19: Append a new `RelationAddress` structure RA_{Ni} to p_N's list
20: Store references to ρ_{iN}, p_i, RA_{Ni} and p_N's list of `RelationAddress`es at RA_{Ni}
21: Store reference to RA_{Ni} at ρ_{iN}
22: **end for**
23: Recompute a using (16)
24: **end function**

`TotalPopulation` is valid since all the propensities of the reacitons involving S_i are zeros and the remaining propensity sums has not changed; it also contains no invalid references.

The final complexity of the deletion operation is $\mathcal{O}(\alpha N)$.

Algorithm 2 Deleting a `Population`

1: **function** DeletePopulation(tp, p_i)
2: **for all** `Relation`s ρ_{ij} in p_i's list **do**
3: Follow ρ_{ij}'s reference to `RelationAddress` pointing at it, RA_{ji}
4: Use the reference at RA_{ji} to itself and to the list holding it to remove it from the list
5: **end for**
6: **for all** `RelationAddress` RA_{ij} in p_i's list **do**
7: Follow RA_{ij}'s reference to the `Relation` ρ_{ji} and to the `Population` p_j owning it
8: Remove ρ_{ji} from p_j's list
9: **end for**
10: Delete all the date in the p_i structure
11: Remove p_i from tp's list
12: **end function**

G. Sampling stage

Similarly to DM[16] and PDM[34], our algorithm simulates the system by randomly sampling time to the next reaction and the reaction itself with certain distributions. Here we describe how it happens in our algorithm.

We begin by generating two random numbers r_1 and r_2. The first one is used to compute the time to the next reaction τ exactly as in DM and PDM (see eq. (3)). The second random number r_2 is used to sample the reaction similarly to how its done in PDM[34].

The sampling process has three stages (see algorithm 3). During the first stage (lines 2-6) we determine the first specie participating in the reaction to happen. To this end, we go through the list of `Population`s $p_1 ... p_N$ until the following condition is satisfied:

$$\sum_{i=1}^{J} \Sigma_i \leqslant a r_2 < \sum_{i=1}^{J+1} \Sigma_i. \tag{17}$$

The second stage (lines 7-11) involves finding the second reactant. We look for a `Relation` ρ_{JK} among those attached to the `Population` p_J for which the following condition holds:

$$\sum_{i=1}^{J} \Sigma_i + n_J \sum_{i=1}^{K} \Psi_{Ji}^{(J)} \leqslant a r_2 < \sum_{i=1}^{J} \Sigma_i + n_J \sum_{i=1}^{K+1} \Psi_{Ji}^{(J)} \tag{18}$$

During the third stage (lines 12-16), we determine which of the reactions possible between the two species is going to happen. We go through ρ_{JK}'s list of the `Reaction`s, looking for a reaction R_{JKL} such that

$$\begin{aligned} \sum_{i=1}^{J} \Sigma_i + n_J \sum_{i=1}^{K} \Psi_{Ji}^{(J)} + n_J \sum_{i=1}^{L} \pi_{JKi}^{(J)} \leqslant a r_2 < \\ < \sum_{i=1}^{J} \Sigma_i + n_J \sum_{i=1}^{K+1} \Psi_{Ji}^{(J)} + n_J \sum_{i=1}^{L+1} \pi_{JKi}^{(J)}. \end{aligned} \tag{19}$$

Note how equations (17) and (18–19) are similar, but their implementation in the pseudocode (algorithm 3) is different. This design minimizes sampling errors due to floating point representation, making the first stage exact.

The resulting reaction sampling finds a reaction in exactly the same manner as equation (4) does. However, the first and the second stages of this sampling process take $\mathcal{O}(N)$ steps and the third step takes $\mathcal{O}(\alpha)$ steps, resulting in a total time complexity of $\mathcal{O}(N + \alpha)$. Using (4) directly requires $\mathcal{O}(M)$ steps[16], which is $\mathcal{O}(\alpha N^2)$ for densely connected reaction networks.

H. Updating stage

When the reaction to occur R_{JKL} is known, our algorithm proceeds to update the data to reflect the changes

Algorithm 3 Reaction sampling

function SAMPLEREACTION(tp, r_2)
  $s_1 \leftarrow 0$, $s_2 \leftarrow 0$
  while $ar_2 > s_1$ **do**
    Get a new `Population` p_i from tp's list
    $s_2 \leftarrow s_1$, $s_1 \leftarrow s_1 + \Sigma_i$
  end while
  $g \leftarrow (ar_2 - s_2)/n_i$
  while $g > 0$ **do**
    Get a new `Relation` ρ_{ij} from p_i's list
    $g \leftarrow g - \Psi_{ij}^{(i)}$
  end while
  $g \leftarrow g + \Psi_{ij}^{(i)}$
  while $g > 0$ **do**
    Get a new `Reaction` R_{ijk} from ρ_{ij}'s list
    $g \leftarrow g - \pi_{ijk}^{(i)}$
  end while
  return R_{ijk}
end function

in species' populations and propensities (see also algorithm 4). For every specie S_i involved in R_{JKL} we read its ID_i from the array stored at R_{JKL}. We run a sequential search for this specie's `Population` p_i over the list kept in `TotalPopulation`. If the specie's `Population` is found, its molecular count n_i is updated using the stoichiomethic coefficient $\sigma_i(R_{JKL})$:

$$n_i \leftarrow n_i + \sigma_i(R_{JKL}). \tag{20}$$

Stoichiometric coefficients are negative for reagents, so their molecular counts may become zero after this step.

After updating the molecular count, we also update all partial and total propensities which depend on it. From every `RelationAddress` RA_{ij} in p_i's list we obtain references to a `Relation` ρ_{ji} in which S_i participates and to the `Population` p_j of its owner specie. For ρ_{ji} we recompute all $\pi_{jik}^{(j)}$ and $\Psi_{ji}^{(j)}$ from scratch using formulas (12) and (13). For p_j, we update the propensity sums as follows:

$$\begin{aligned} \Lambda_j^{(j)} &\leftarrow \Lambda_j^{(j)} - \Psi_{ji}^{(j)}, \\ \Sigma_j &\leftarrow n_j \Lambda_j^{(j)}. \end{aligned} \tag{21}$$

Since each of the species involved in R_{JKL} may be involved in $\mathcal{O}(\alpha N)$ reaction, updating the structure in this way takes a total of $\mathcal{O}(\alpha N)$ operations.

If the specie S_i is a product, its `Population` may not exist yet. In this case we add a new `Population` of S_i using its ID_i as described in section III D. The molecular count of the newly added specie is its stoichiometric coefficient in R_{JKL}, $\sigma_i(R_{JKL})$. Additions take $\mathcal{O}(\alpha N)$ operations.

When we're done updating the existing `Populations` and adding the new ones, we iterate through the list of `Populations` again and delete the ones with a molecular count of zero as described in section III F. The deletion takes $\mathcal{O}(\alpha N)$.

Finally, we recompute the total propensity of the system a using (16).

Algorithm 4 Post-sampling data update

function UPDATEEXISTINGPOPULATIONS(tp, R_{IJK})
  for all species S_i participating in R_{IJK} **do**
    Search for a `Population` of specie S_i in tp's list using its ID_i
    if S_i has a `Population` p_i in tp's list **then**
      $n_i \leftarrow n_i + \sigma_i(R_{IJK})$
      for all `RelationAddress`es RA_{ij} in p_i's list **do**
        Follow the references at RA_{ij} to get the `Relation` ρ_{ji} and its owner `Population` p_j
        $\Lambda_j^{(j)} \leftarrow \Lambda_j^{(j)} - \Psi_{ji}^{(j)}$
        $\Psi_{ji}^{(j)} \leftarrow 0$
        for all `Reaction`s R_{jik} in ρ_{ji}'s list **do**
          Recompute $\pi_{jik}^{(j)}$ using (12)
          $\Psi_{ji}^{(j)} \leftarrow \Psi_{ji}^{(j)} + \pi_{jik}^{(j)}$
        end for
        $\Lambda_j^{(j)} \leftarrow \Lambda_j^{(j)} + \Psi_{ji}^{(j)}$
        $\Sigma_j \leftarrow n_j \cdot \Lambda_j^{(j)}$
      end for
    else
      ADDPOPULATION(tp, ID_i, $\sigma_i(R_{IJK})$)
    end if
  end for
  for all `Population`s p_i in tp's list **do**
    if $n_i == 0$ **then**
      DELETEPOPULATION(tp, p_i)
    end if
  end for
end function

I. Summary of EPDM

EPDM is an SSA which only maintains the data about the species with nonzero molecular count (see algorithm 5). This ensures that the number of tracked molecular species N is as low as possible.

Our algorithm uses a data structure described in section III C. The structure holds one entry for each possible reaction, bringing storage requirements of our algorithm to $\mathcal{O}(\alpha N^2)$.

The data structure is initialized by constructing it to be empty, then adding the specie data for every specie initially present in the system. The process is described in section III E and takes $\mathcal{O}(\alpha N^2)$ operations.

Each step of the simulation executes a single reaction. It is composed of a sampling step (see section III G) and an updating step (section III H). Each of these takes $\alpha\mathcal{N}$ operations, so the total number of operations needed to simulate one reaction is also $\mathcal{O}(\alpha N)$.

Algorithm 5 EPDM overview

```
▷ Initialization stage, see section III E
Create TotalPopulation tp based on initial conditions
while a > 0 and termination conditions not met do
    ▷ Sampling stage, see section III G
    Generate two random numbers r1 and r2
    Compute τ using (3)
    R ←SampleReaction(tp, r2)
    ▷ Updating stage, see section III H
    for all species Si involved in R do
        Search tp's list for Population of Si
        if a Population pi of Si has been found then
            Update Si's molecular count ni
            Update all partial propensities which
              depend on ni
        else if no Population of Si has been found then
            Add a Population of Si to tp's list
            ▷ details in section III D
        end if
    end for
    Find all Populations with n == 0 in tp's list
    Remove all found Populations
    ▷ details in section III F
end while
```

When some reaction produces any number of molecules of a previously unknown specie, specie data is added as described in section III D. To generate the reactions dynamically, a user-defined function `reactions()` is used which takes two species and produces a list of reactions between them, complete with rates.

When any specie's molecular count reaches zero, its data is pruned from the structure as described in section III F.

J. Implementation

We implemented our algorithm as a C++ framework. The user must define a class `Specie` with a constructor from a string ID which must save the string into the member variable `m_id`. The class must define a method `Specie::reactions(Specie)` returning a list of possible `Reaction`s between the caller `Specie` and the argument.

After implementing the class `Specie`, users can simulate the system. Two stopping criteria are currently available by default: the algorithm can stop either when a certain number of reactions have been executed or a certain simulation time has passed.

A global dictionary with arbitrary parameters loaded from a configuration file is provided for convenience.

The implementation is currently available for Linux and Mac OS X. It is tested with GNU gcc 4.9.3 and GNU make 4.1.

The code is available at `https://github.com/abernatskiy/epdm`.

IV. BENCHMARKS

We benchmark performance of EPDM against two direct methods: PDM[343] and DM[384]

Because our model was designed with complex systems in mind we studied performance only for strongly coupled systems. In both systems every specie can interact with every other specie in a unique way, ensuring that the number of reactions is

$$M = \frac{N(N+1)}{2} \tag{22}$$

and

$$\alpha = 1. \tag{23}$$

Both models are designed in such a way that the total number of species is preserved throughout the simulation time.

Models below are designed to measure the performance of our method. They don't fully illustrate the power of the model because they have fixed number molecules, which is necessary to keep for benchmark. The most striking performance gain is achieved when listing all the species isn't possible in principle or due to computational costs. For example, in case of realistic polymerization and autocatalysis model used to study prebiotic polymerization Guseva et al. it was possible to increase the maximum length of simulated polymers from 12 to 25 by employing our algorithm. Note that limit of 25 wasn't due to restrictions of our algorithm, but due to necessity to calculate minimum energy folding configuration of every chain, which is an NP-complete problem.

A. CPU time of EPDM is linear for a strongly coupled system

The first model ("colliding particles") is made to test how our algorithm performs on chemical systems where no new molecules are created and no molecules ever disappear. This is a model of a system consisting of colliding particles of N species. Particles behave like rigid spheres: they collide, and bounce back without internal changes. All of the particle species are known in advance, none are added or removed over the course of the simulation. The system is defined by the following set of equations:

$$P_i + P_j \xrightarrow{k} P_i + P_j, \quad i, j \in [1, N]. \tag{24}$$

In our simulations we vary number of species N from 10 to 7000. Every specie has a population of 50 molecules.

[3] We took implementation from http://mosaic.mpi-cbg.de/pSSALib/pSSAlib.html

[4] We took implementation from http://sourceforge.net/projects/stochkit/

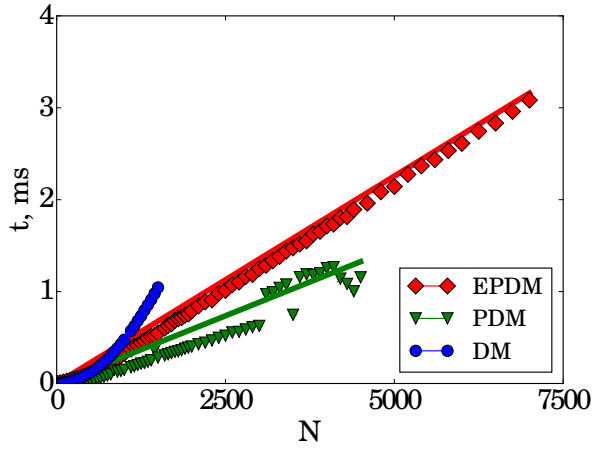

FIG. 2. CPU time per reaction as a function of the number of species in the system for "colliding particles" model (section IV A) for EPDM, PDM and DM.

Collision rate k is fixed at 0.5 s^{-1}. Every simulation runs until 5000 reactions have occurred. For every value of N, CPU time to simulation completion was measured 10 times and the average time is reported.

Figure 2 shows CPU time it takes per reaction for DM, PDM and EPDM. DM is clearly quadratic. For any given N PDM outperforms the EPDM, but both are linear in time. It is important to note that DM and PDM were stopped for a relatively low values of N due to excessive RAM consumption (about ~ 120GB) by both of the applications, which we suspect was due to implementation issues (in particular, in XML handling libraries).

B. CPU time stays linear when species are actively deleted and created

The second test checks if algorithm keeps linear time when `Populations` are added and removed from the system. The test system ("particles with color") consists of N colliding particles of N types, each of which has an internal property ("color") that is changed in the collision. The following equation defines the system:

$$P_i^\alpha + P_j^\beta \xrightarrow{k} P_i^{\alpha+1} + P_j^{\beta+1}, \quad i,j \in [1,N],\ \alpha,\beta \in [1,\Omega] \tag{25}$$

Indexes i, j enumerate particle types and run from 1 to N. Indexes α, β enumerate colors of particles; particle can have one of Ω colors. During the collision color index of a participating particle goes up $P_i^\alpha \to P_i^{\alpha+1}$, until it reaches maximal index Ω, after which it drops back to 0: $P_i^\Omega \to P_i^0$.

Since every combination of a particle type and a color is considered a separate specie, every reaction causes two species to go extinct and their `Populations` to be deleted. It also adds two new species, which requires adding two new `Populations`. Thus, the total number of species simultaneously present in the system is maintained at exactly the same level. Every specie is represented in the system by a single molecule.

To run such a simulation in DM and PDM frameworks we had to enumerate all $N\Omega$ of the possible species. This slows down the simulation enough to make the comparison impossible beyond a small number of species. The figure 3 shows how CPU time per reaction depends on the number of species in EPDM.

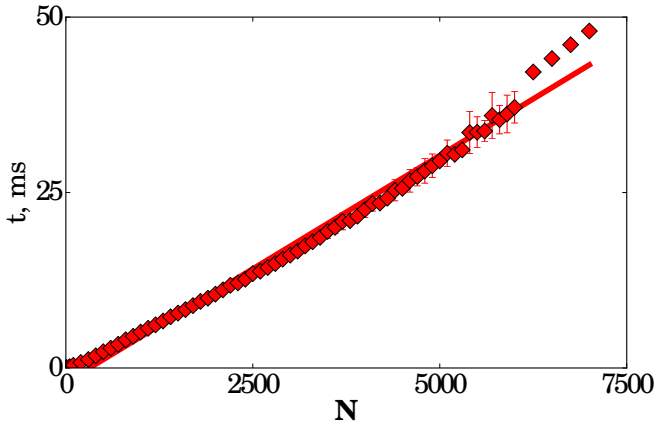

FIG. 3. CPU time per reaction for EPDM simulation of "particles with color" (see section IV B).

V. CONCLUSION AND DISCUSSION

Stochastic simulations are actively used in molecular and systems biology. The bigger and more complex the system, the more important the performance of the simulation algorithm becomes. It is also more burdensome or even impossible to list all the species and reactions for complex systems. We introduced a general purpose, exact stochastic simulation algorithm which allows to avoid listing all the possible species and reactions by defining the general rules governing the system instead. Built within the partial propensity framework[34,35], our algorithm achieves linear time complexity in the number of molecular species.

The algorithm has its limitations. First, it is limited to reactions of maximum of two reactants. In the chemical system that doesn't present much of an issue because reactions with three molecules are significantly rarer then binary reactions and more complex reactions can be represented as sequences of binary reactions. Second, it cannot simulate spatially nonhomogeneous systems. However, as long as all reactions involve no more than two reactants and the observation from section II A holds, it is possible to simulate any system for which the set of reactions between any two species is known.

Benchmarks suggest that our algorithm is slower than PDM by a constant factor. Therefore, one should use

PDM when it is unlikely that a significant proportion of species will have a molecular count of zero.

Our results suggest that in complex systems (such as polymerization reactions and networks of intra-cellular reactions) and in the case of emergent phenomena studies (e.g. origins of life) EPDM can significantly improve the performance of the stochastic simulations. The software implementation of the algorithm is available as an open source public repository `https://github.com/abernatskiy/epdm`.

VI. ACKNOWLEDGMENTS

Authors appreciate support from the NSF grant MCB1344230.

[1]Arkin, A., Ross, J., and McAdams, H. H. (1998). Stochastic kinetic analysis of developmental pathway bifurcation in phage lambda-infected Escherichia coli cells. *Genetics*, 149(4):1633–48.

[2]Auger, A., Chatelain, P., and Koumoutsakos, P. (2006). R-leaping: Accelerating the stochastic simulation algorithm by reaction leaps. *The Journal of Chemical Physics*, 125(8):084103.

[3]Berman, S., Halasz, A., Hsieh, M., and Kumar, V. (2009). Optimized Stochastic Policies for Task Allocation in Swarms of Robots. *IEEE Transactions on Robotics*, 25(4):927–937.

[4]Berman, S., Halasz, A., Kumar, V., and Pratt, S. (2007). Bio-Inspired Group Behaviors for the Deployment of a Swarm of Robots to Multiple Destinations. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 2318–2323. IEEE.

[5]Breban, R., Drake, J. M., Stallknecht, D. E., and Rohani, P. (2009). The Role of Environmental Transmission in Recurrent Avian Influenza Epidemics. *PLoS Computational Biology*, 5(4):e1000346.

[6]Cai, X. and Xu, Z. (2007). K -leap method for accelerating stochastic simulation of coupled chemical reactions. *Journal of Chemical Physics*, 126(7):1–10.

[7]Cao, Y., Gillespie, D. T., and Petzold, L. R. (2005a). Avoiding negative populations in explicit Poisson tau-leaping. *The Journal of Chemical Physics*, 123(5):054104.

[8]Cao, Y., Gillespie, D. T., and Petzold, L. R. (2005b). The slow-scale stochastic simulation algorithm. *The Journal of Chemical Physics*, 122(1):014116.

[9]Cao, Y., Gillespie, D. T., and Petzold, L. R. (2006). Efficient step size selection for the tau-leaping simulation method. *The Journal of Chemical Physics*, 124(4):044109.

[10]Cao, Y., Li, H., and Petzold, L. (2004). Efficient formulation of the stochastic simulation algorithm for chemically reacting systems. *The Journal of Chemical Physics*, 121(9):4059.

[11]Caravagna, G., dOnofrio, A., Milazzo, P., and Barbuti, R. (2010). Tumour suppression by immune system through stochastic oscillations. *Journal of Theoretical Biology*, 265(3):336–345.

[12]Carbone, G. and Giannoccaro, I. (2015). Model of human collective decision-making in complex environments. *The European Physical Journal B*, 88(12):339.

[13]Dieckmann, U. and Law, R. (1996). The dynamical theory of coevolution: a derivation from stochastic ecological processes. *Journal of Mathematical Biology*, 34(5-6):579–612.

[14]Doob, J. L. (1953). *Stochastic Processes*. John Wiley & Sons, Inc.;Chapman & Hall, New York, New York, USA.

[15]Gibson, M. A. and Bruck, J. (2000). Efficient Exact Stochastic Simulation of Chemical Systems with Many Species and Many Channels. *The Journal of Physical Chemistry A*, 104(9):1876–1889.

[16]Gillespie, D. T. (1976). A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434.

[17]Gillespie, D. T. (1977). Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361.

[18]Gillespie, D. T. (1992). A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(1-3):404–425.

[19]Gillespie, D. T. (2001). Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716.

[20]Guseva, E. A., Zuckermann, R. N., and Dill, K. A. (2016). How did prebiotic polymers become informational foldamers?

[21]Henkelman, G. and Jonsson, H. (2001). Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table. *The Journal of Chemical Physics*, 115(21):9657.

[22]Hillston, J. (2005). Fluid flow approximation of PEPA models. In *Second International Conference on the Quantitative Evaluation of Systems (QEST'05)*, pages 33–42. IEEE.

[23]Juanico, D. E. (2012). Critical network effect induces business oscillations in multi-level marketing systems.

[24]Karlebach, G. and Shamir, R. (2008). Modelling and analysis of gene regulatory networks. *Nature Reviews Molecular Cell Biology*, 9(10):770–780.

[25]Keeling, M. J. and Rohani, P. (2008). *Modeling infectious diseases in humans and animals*. Princeton University Press.

[26]Le Galliard, J.-F., Ferrière, R., and Dieckmann, U. (2005). Adaptive evolution of social traits: origin, trajectories, and correlations of altruism and mobility. *The American naturalist*, 165(2):206–24.

[27]Legrand, J., Grais, R. F., Boelle, P. Y., Valleron, A. J., and Flahault, A. (2007). Understanding the dynamics of Ebola epidemics. *Epidemiology and infection*, 135(4):610–21.

[28]McCollum, J. M., Peterson, G. D., Cox, C. D., Simpson, M. L., and Samatova, N. F. (2006). The sorting direct method for stochastic simulation of biochemical systems with varying reaction execution behavior. *Computational Biology and Chemistry*, 30(1):39–49.

[29]McKane, A. J. and Newman, T. J. (2005). Predator-prey cycles from resonant amplification of demographic stochasticity. *Physical review letters*, 94(21):218102.

[30]Mobilia, M. (2012). Stochastic dynamics of the prisoner's dilemma with cooperation facilitators. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 86(1 Pt 1):011134.

[31]Montero, M. (2008). Predator-Prey Model for Stock Market Fluctuations. *SSRN Electronic Journal*.

[32]Peng, X., Zhou, W., and Wang, Y. (2007). Efficient binomial leap method for simulating chemical kinetics. *The Journal of Chemical Physics*, 126(22):224109.

[33]Peng, X.-j. and Wang, Y.-f. (2007). L-leap: accelerating the stochastic simulation of chemically reacting systems. *Applied Mathematics and Mechanics*, 28(10):1361–1371.

[34]Ramaswamy, R., González-Segredo, N., and Sbalzarini, I. F. (2009). A new class of highly efficient exact stochastic simulation algorithms for chemical reaction networks. *The Journal of chemical physics*, 130(24):244104.

[35]Ramaswamy, R. and Sbalzarini, I. F. (2010). A partial-propensity variant of the composition-rejection stochastic simulation algorithm for chemical reaction networks. *Journal of Chemical Physics*, 132(4):1–6.

[36]Rathinam, M., Petzold, L. R., Cao, Y., and Gillespie, D. T. (2003). Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. *The Journal of Chemical Physics*, 119(24):12784.

[37]Reichenbach, T., Mobilia, M., and Frey, E. (2007). Mobility promotes and jeopardizes biodiversity in rockpaperscissors games. *Nature*, 448(7157):1046–1049.

[38]Sanft, K. R., Wu, S., Roh, M., Fu, J., Lim, R. K., and Petzold, L. R. (2011). StochKit2: software for discrete stochastic simulation of biochemical systems with events. *Bioinformatics (Oxford, England)*, 27(17):2457–8.

[39]Shkarayev, M. S., Schwartz, I. B., and Shaw, L. B. (2013). Recruitment dynamics in adaptive social networks. *Journal of physics. A, Mathematical and theoretical*, 46(24):245003.

[40]Wilkinson, D. J. (2011). *Stochastic Modelling for Systems Biology*. CRC Press.